Lower resolution X-ray spectroscopy

Keith Arnaud

NASA Goddard

University of Maryland

Practical X-ray spectroscopy

Most X-ray spectra are of moderate or low resolution (eg Chandra ACIS or XMM-Newton EPIC).

However, the spectra generally cover a bandpass of more than 1.5 decades in energy.

Moreover, the continuum shape often provides important physical information.

Therefore, unlike in the optical, most uses of X-ray spectra have involved a simultaneous analysis of the entire spectrum rather than an attempt to measure individual line strengths.

Martin Elvis

3C 273

Optical Spectrum

Proportional counter
e.g. ROSAT PSPC

CCD
e.g. Chandra ACIS

Grating

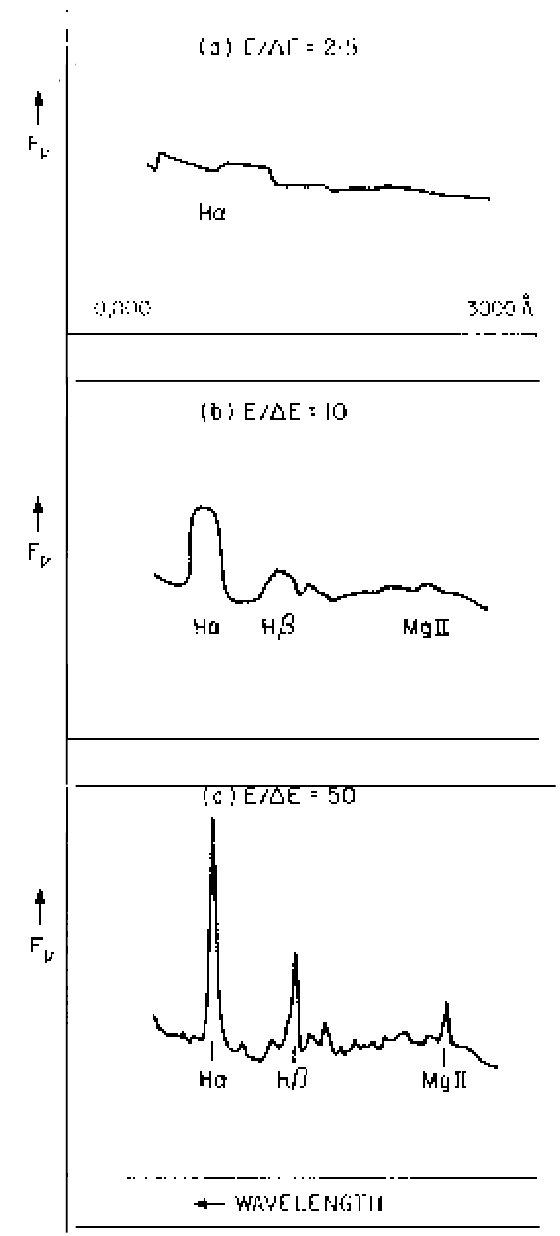

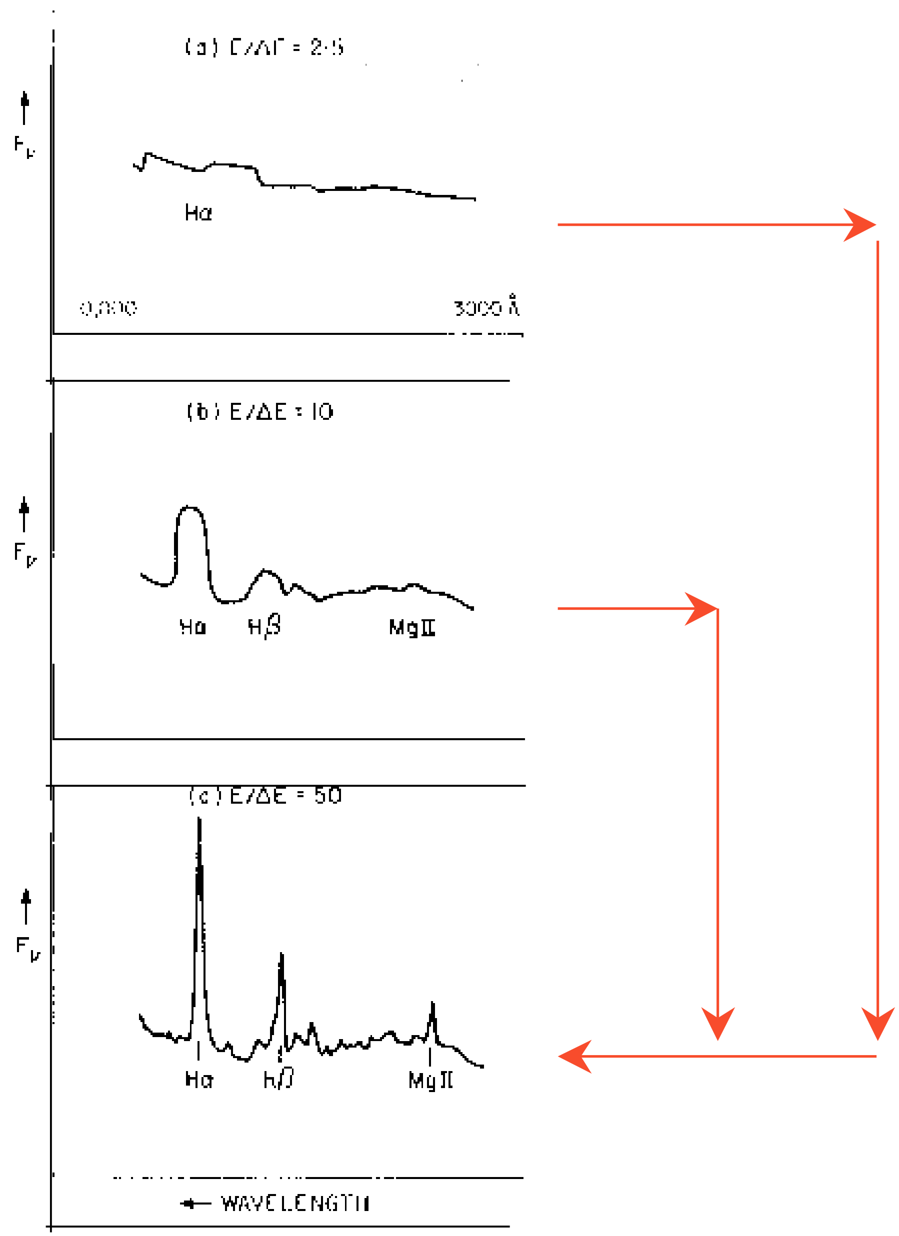

Can we start
with these…

and deduce this ?

Can we start with this…

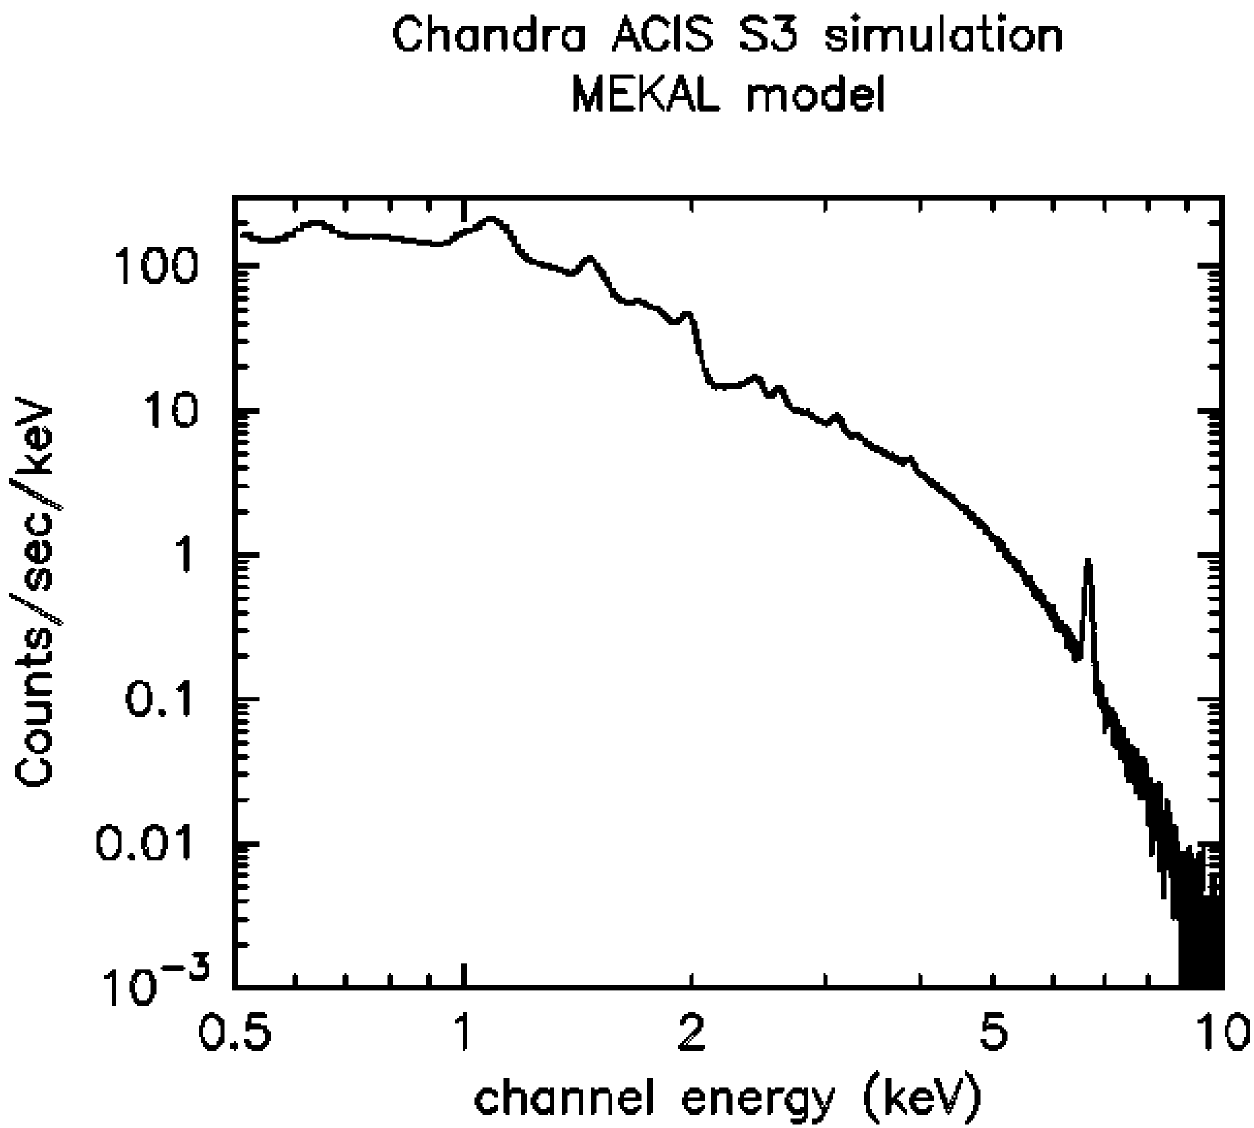

and deduce this

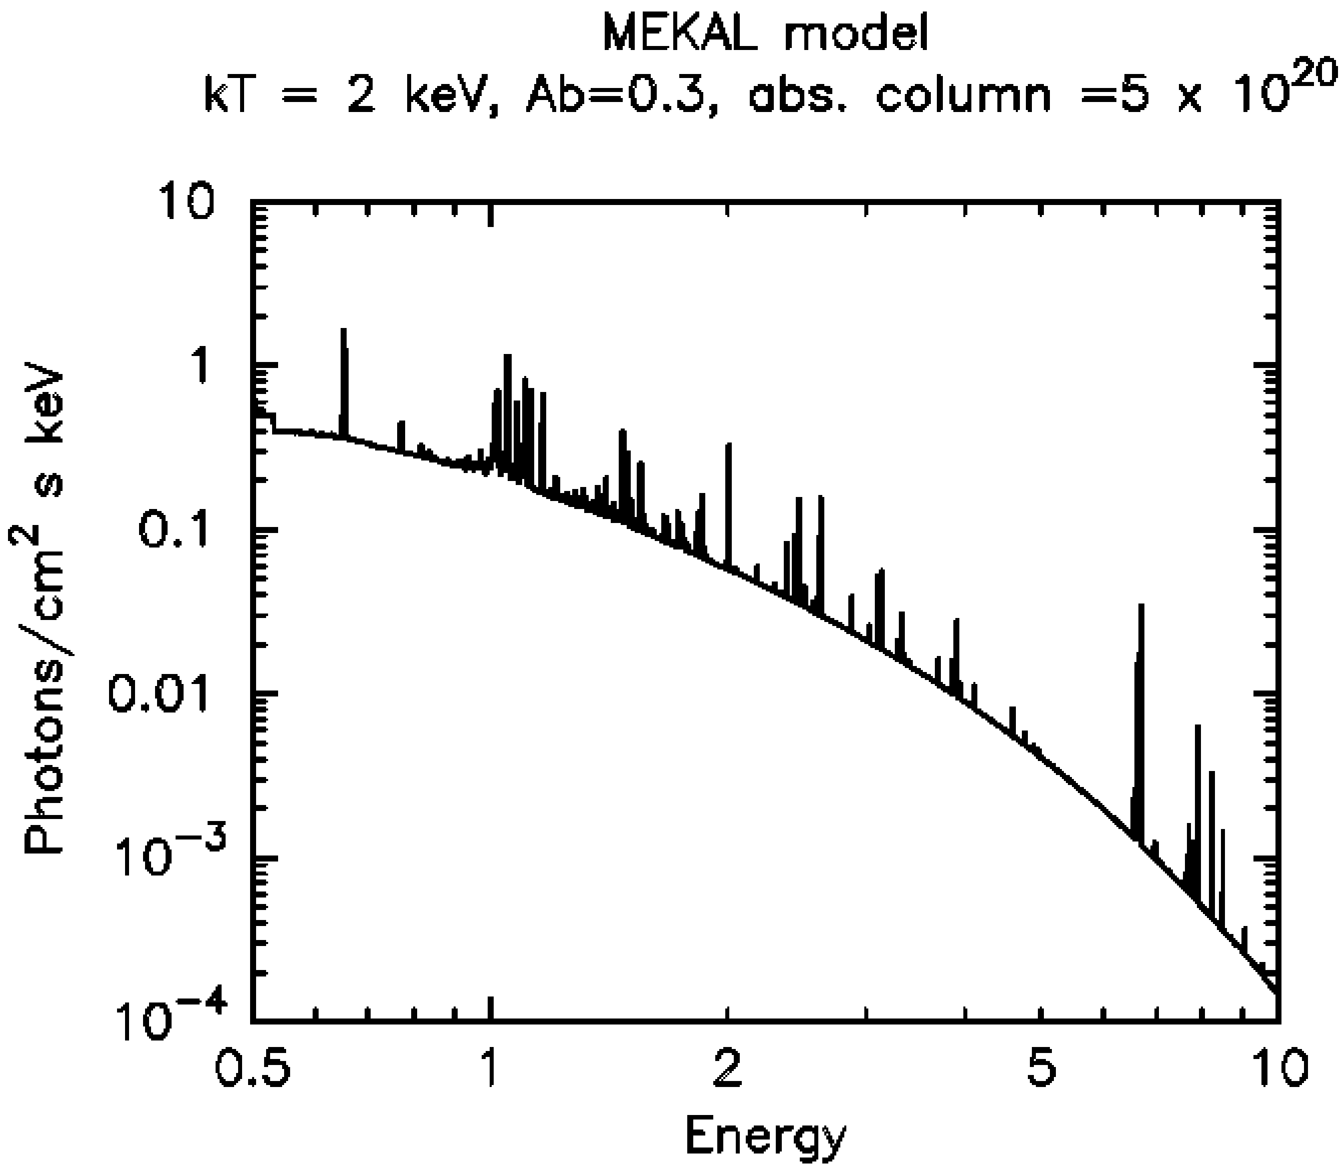

The Basic Problem

Suppose we observe $D(I)$ counts in channel I (of N) from some source. Then :

$$D(I) = T \int R(I,E)\, A(E)\, S(E)\, dE$$

- T is the observation length (in seconds)
- $R(I,E)$ is the probability of an incoming photon of energy E being registered in channel I (dimensionless)
- $A(E)$ is the energy-dependent effective area of the telescope and detector system (in cm^2)
- $S(E)$ is the source flux at the front of the telescope (in $photons/cm^2/s/keV$

An example R(I,E)

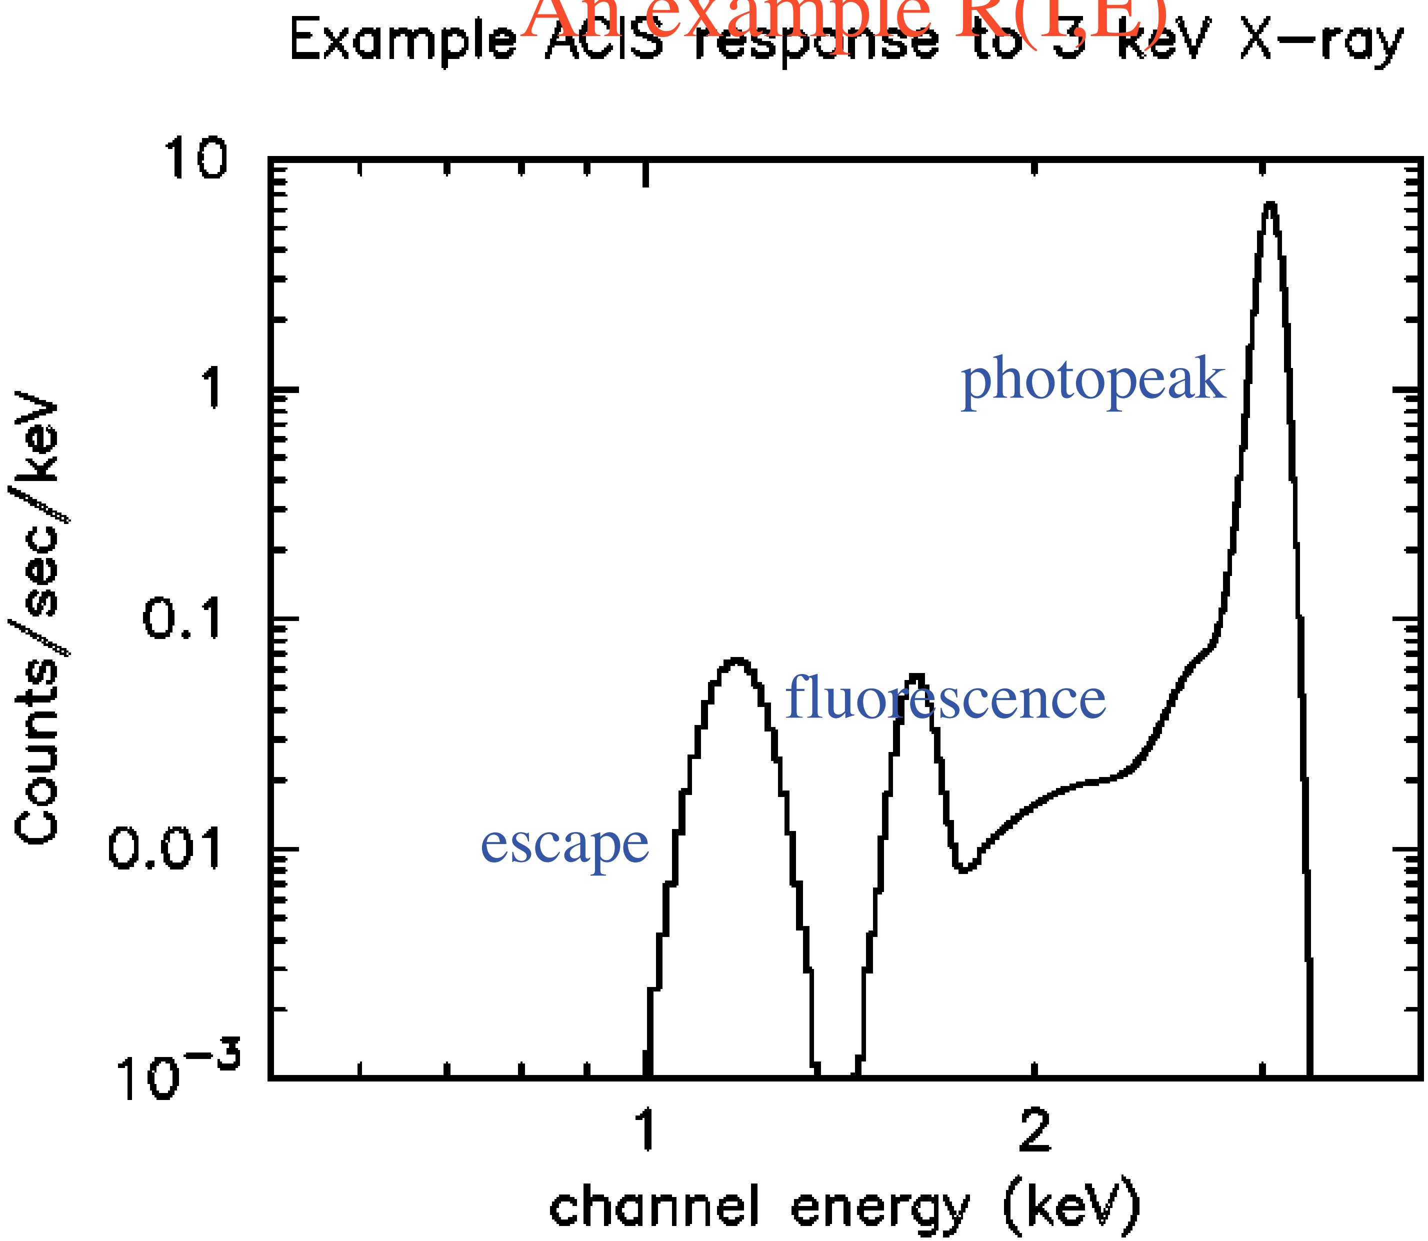

Example A(E)

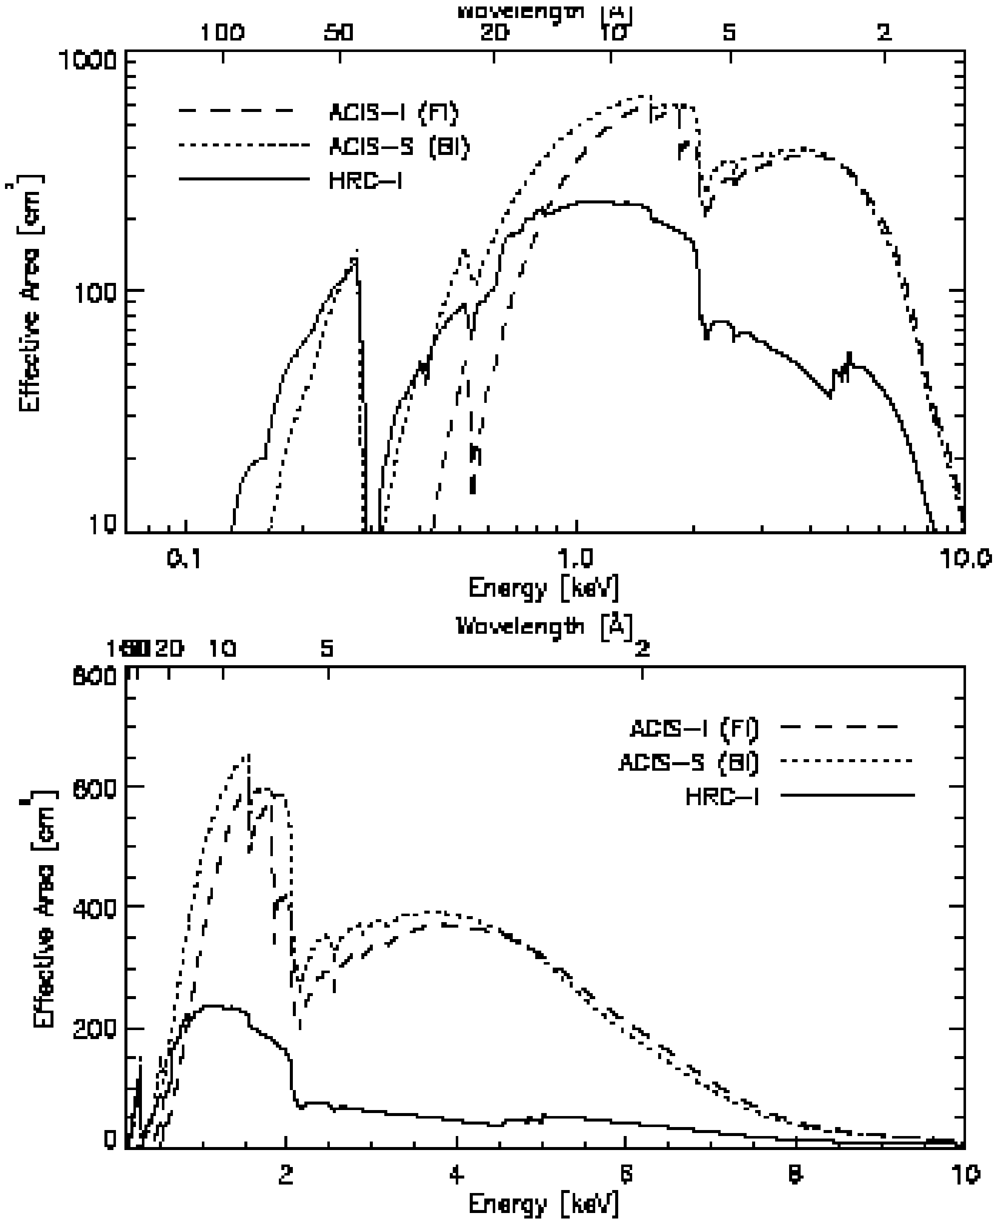

Wavelength [Å]
100
50
20
10
5
2
1000
100
10
ACIS–I (FI)
ACIS–S (BI)
HRC–I
Effective Area [cm²]
0.1
1.0
10.0
Energy [keV]
Wavelength [Å]
20
10
5
2
800
600
400
200
0
ACIS–I (FI)
ACIS–S (BI)
HRC–I
Effective Area [cm²]
2
4
6
8
10
Energy [keV]

Analogs in optical/UV

An example $R(I,E)$ would be the resolution of a spectrometer. In most optical/UV instruments $R(I,E)$ is simple - a Gaussian or Lorentzian shape.

$A(E)$ is the product of telescope reflectivities, detector efficiencies, and filter transmissions. In optical/UV astronomy we usually divide the observed data by this function to obtain the fluxed spectrum.

But can I ignore the response ?

Sometimes, yes.

If you have CCD (eg Chandra ACIS or XMM-Newton EPIC) spectra you must use the response, R(I,E).

If you have Chandra HETG spectra then you can treat them like optical/UV spectra.

However, for XMM-Newton RGS if you try to do this you will get incorrect results. The RGS spectral response has wide wings so line fluxes will be wrong and the continuum level overestimated.

The Basic Problem II

$$D(I) = T \int R(I,E)\, A(E)\, S(E)\, dE$$

We assume that T, A(E) and R(I,E) are known and want to solve this integral equation for S(E). We can divide the energy range of interest into M bins and turn this into a matrix equation :

$$D_i = T \sum R_{ij} A_j S_j$$

where S_j is now the flux in photons/cm^2/s in energy bin J. We want to find S_j.

The Basic Problem III

$$D_i = T \sum R_{ij} A_j S_j$$

The obvious tempting solution is to calculate the inverse of R_{ij}, premultiply both sides and rearrange :

$$(1/T A_j) \sum (R_{ij})^{-1} D_i = S_j$$

This does not work ! The S_j derived in this way are very sensitive to slight changes in the data D_i. This is a great method for amplifying noise.

A (brief) Mathematical Digression

This should not have come as a surprise to anyone with any data analysis experience. This is the "remote sensing problem" and arises in many areas of astronomy as well as eg geophysics and medical imaging.

In mathematics the integral is known as a Fredholm equation of the first kind. Tikhonov showed that such equations can be solved using "regularization" - applying prior knowledge to damp the noise.

A familiar example is maximum entropy but there are a host of others. Some of these have been tried on X-ray spectra - none have had any impact on the field.

Forward-fitting

The standard method of analyzing X-ray spectra is "forward-fitting". This comprises the following steps…

- Calculate a model spectrum.
- Multiply the result by an instrumental response matrix (R(I,E)*A(E)).
- Compare the result with the actual observed data by calculating some statistic.
- Modify the model spectrum and repeat till the best value of the statistic is obtained.

Forward-fitting algorithm

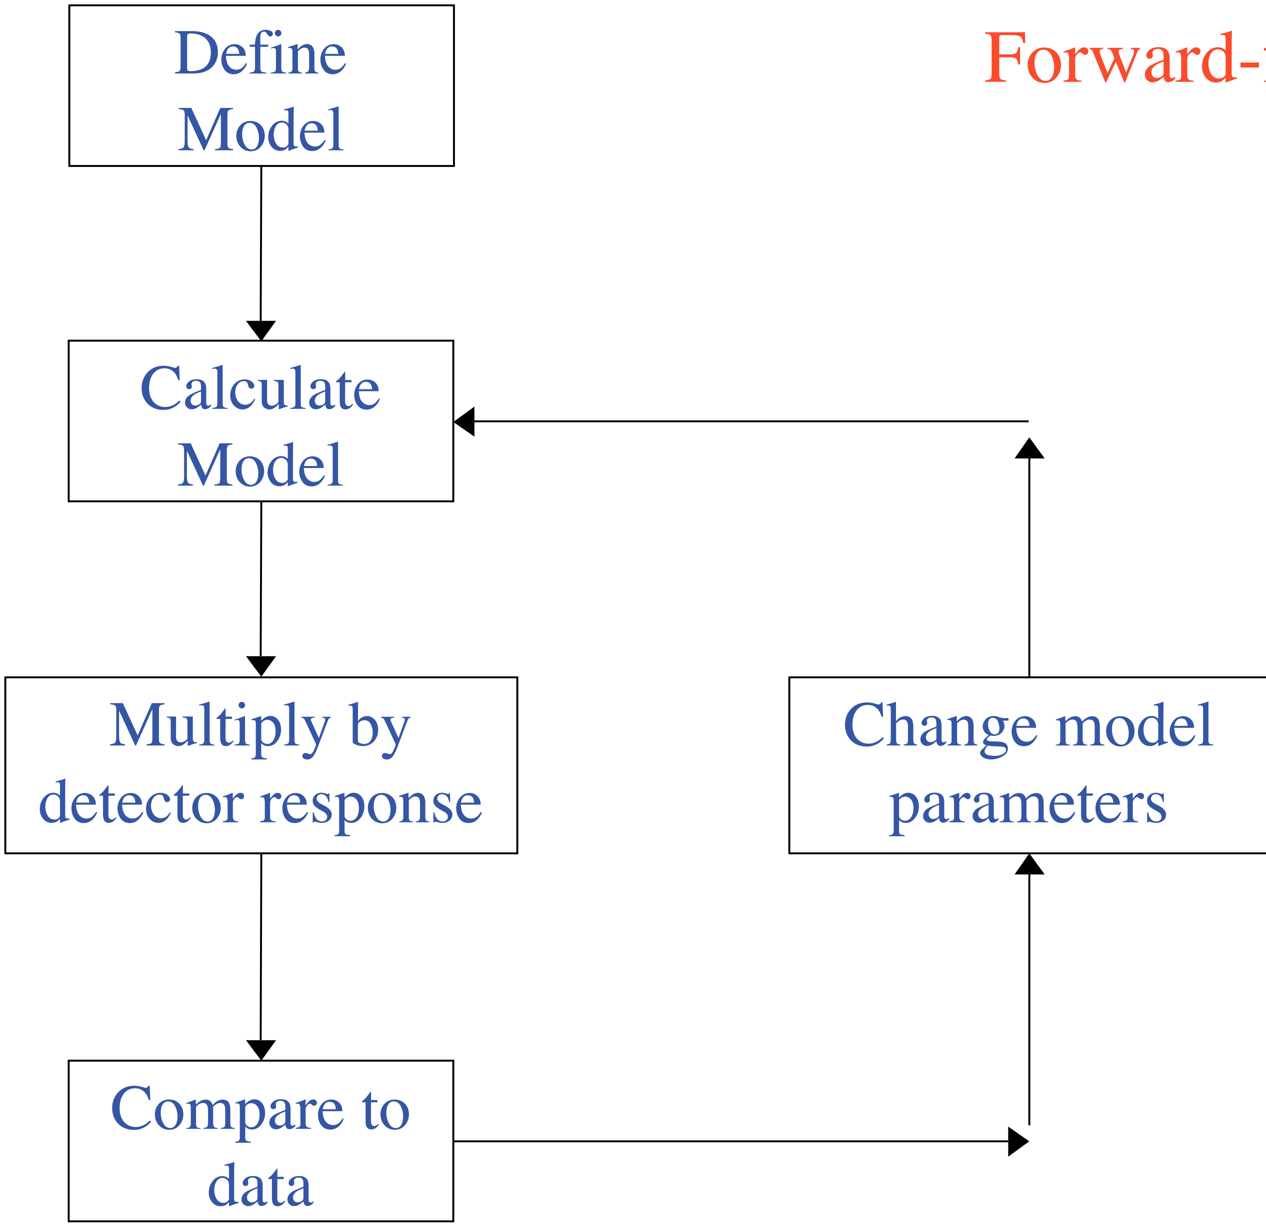

This only works if the model spectrum can be expressed in a reasonably small number of parameters (although I have seen people fit spectra using models with over 100 parameters).

The aim of the forward-fitting is then to obtain the best-fit and confidence ranges of these parameters.

Spectral fitting programs

o XSPEC - part of HEAsoft. General spectral fitting program with many models available.

o Sherpa - part of CIAO. Multi-dimensional fitting program which includes the XSPEC model library and can be used for spectral fitting.

o SPEX - from SRON in the Netherlands. Spectral fitting program specialising in collisional plasmas and high resolution spectroscopy.

o ISIS - from the MIT Chandra HETG group. Mainly intended for the analysis of grating data. Incorporated in Sherpa as GUIDE.

Models

All models are wrong, but some are useful - George Box

X-ray spectroscopic models are usually built up from individual components. These can be thought of as two basic types -additive (an emission component e.g. blackbody, line,…) or multiplicative (something which modifies the spectrum e.g. absorption).

$$\text{Model} = M_1 * M_2 * (A_1 + A_2 + M_3 * A_3) + A_4$$

```
XSPEC>model ?

 Possible additive models are :
 apec     bbody    bbodyrad bexrav   bexriv   bknpower bkn2pow  bmc
 bremss   c6mekl   c6pmekl  c6pvmkl  c6vmekl  cemekl   cevmkl   cflow
 compbb   compLS   compST   compTT   cutoffpl disk     diskbb   diskline
 diskm    disko    diskpn   equil    gaussian gnei     grad     grbm
 laor     lorentz  meka     mekal    mkcflow  nei      npshock  nteea
 pegpwrlw pexrav   pexriv   photoion plcabs   powerlaw posm     pshock
 raymond  redge    refsch   sedov    srcut    sresc    step     vapec
 vbremss  vequil   vgnei    vmeka    vmekal   vmcflow  vnei     vnpshock
 vpshock  vraymond vsedov   zbbody   zbremss  zgauss   zpowerlw atable


 Possible multiplicative models are :
 absori   acisabs  constant cabs     cyclabs  dust     edge     expabs
 expfac   highecut hrefl    notch    pcfabs   phabs    plabs    pwab
 redden   smedge   spline   SSS_ice  TBabs    TBgrain  TBvarabs uvred
 varabs   vphabs   wabs     wndabs   xion     zedge    zhighect zpcfabs
 zphabs   zTBabs   zvarabs  zvfeabs  zvphabs  zwabs    zwndabs  mtable
 etable

 Possible mixing models are :
 ascac    projct   xmmc

 Possible convolution models are :
 gsmooth  lsmooth  reflect  rgsxsrc

 Possible pile-up models are :
 pileup
```

Additive Models

Basic additive (emission) models include :

- blackbody
- thermal bremsstrahlung
- power-law
- collisional plasma (raymond, mekal, apec)
- Gaussian or Lorentzian lines

There are many more models available covering specialised topics such as accretion disks, comptonized plasmas, non-equilibrium ionization plasmas, multi-temperature collisional plasmas…

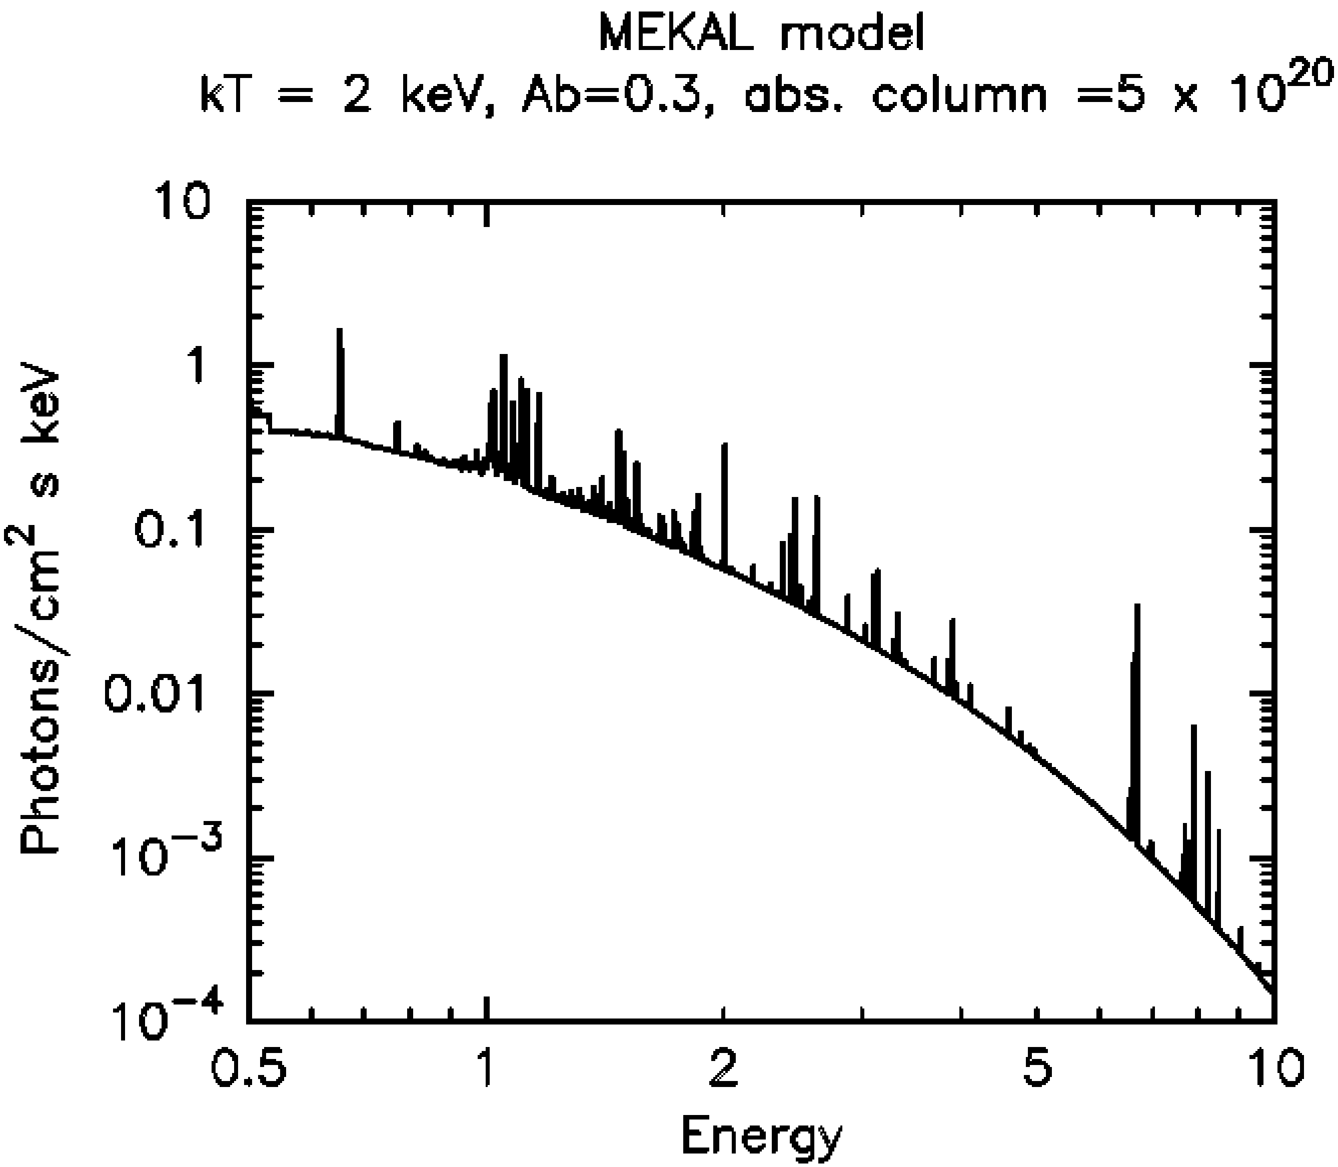
MEKAL model
kT = 2 keV, Ab=0.3, abs. column =5 x 10^20
Photons/cm^2 s keV
10
1
0.1
0.01
10^-3
10^-4
0.5
1
2
5
10
Energy

Multiplicative Models

and multiplicative models include :

- photoelectric absorption due to our Galaxy
- photoelectric absorption due to ionized material
- high energy exponential roll-off.
- cyclotron absorption lines.

Galactic absorption

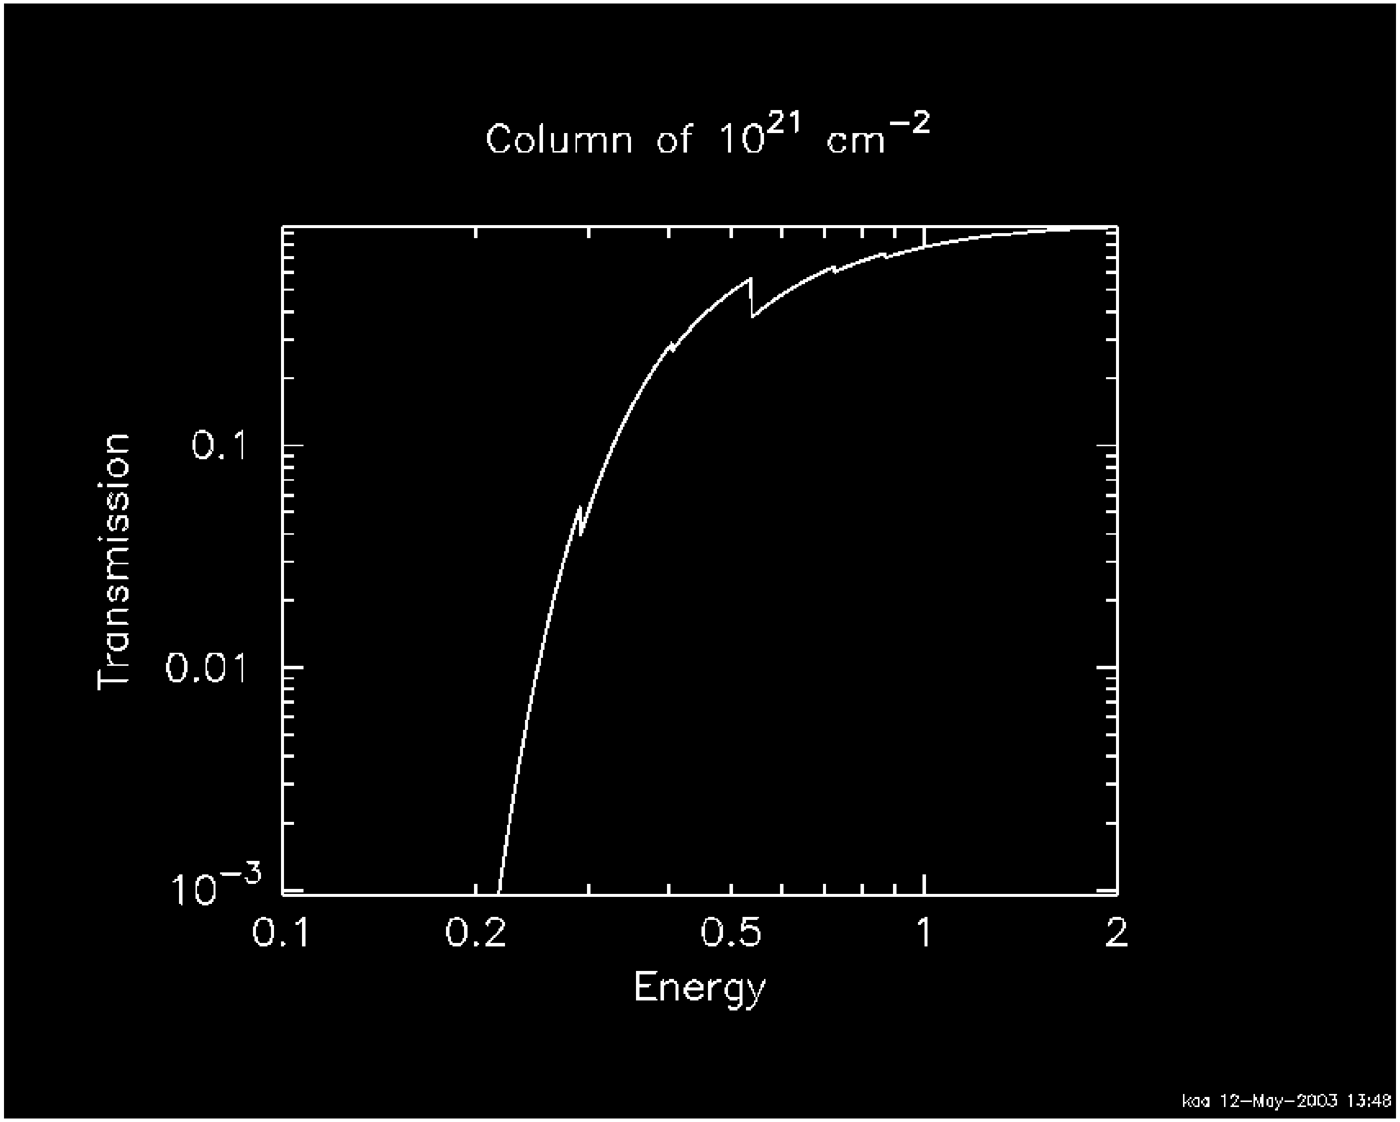

Convolution Models
(for the aficionados)

These are models which take as input the current model and manipulate it in some way. Examples are :

- Smoothing with a Gaussian or Lorentzian function (e.g. velocity broadening)
- Compton reflection
- Pile-up

Roll Your Own Models

There is a simple XSPEC model interface which enables astronomers to write new models and fit them to their data. You can write your own subroutine (in Fortran or C) and hook it in - the subroutine takes in the energies on which to calculate the model and writes out the fluxes (in photons/cm^2/s).

In addition, there is also a standard format for files containing model spectra so these too can be fit to data without having to add new routines to XSPEC.

Finding the best-fit

Finding the best-fit means minimizing the statistic value. There are many algorithms available to do this in a computationally efficient fashion (see Numerical Recipes).

Most methods used to find the best-fit are local i.e. they use some information around the current parameters to guess a new set of parameters. All these methods are liable to get stuck in a local minimum. Watch out for this !

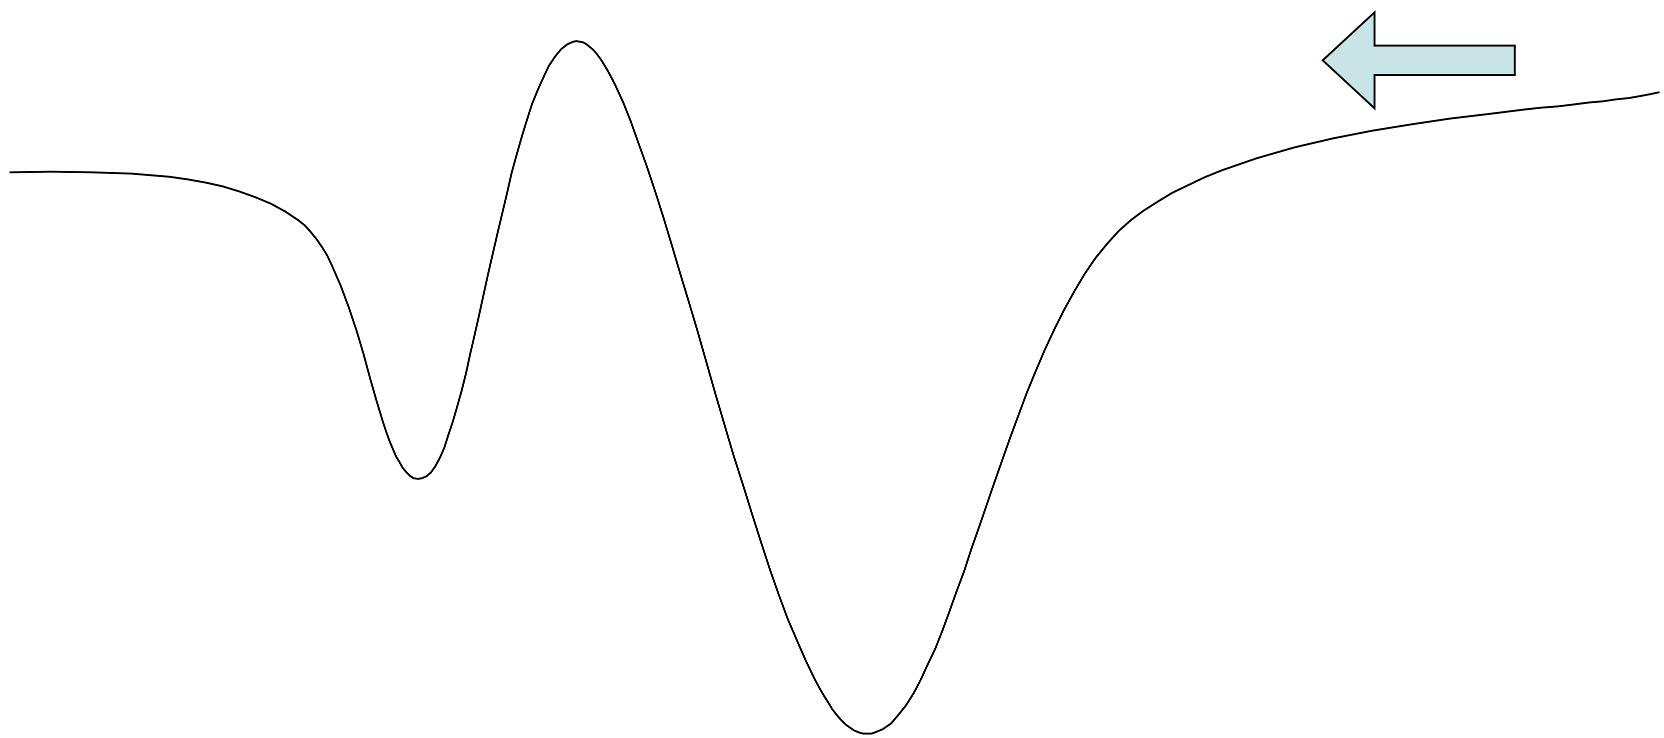

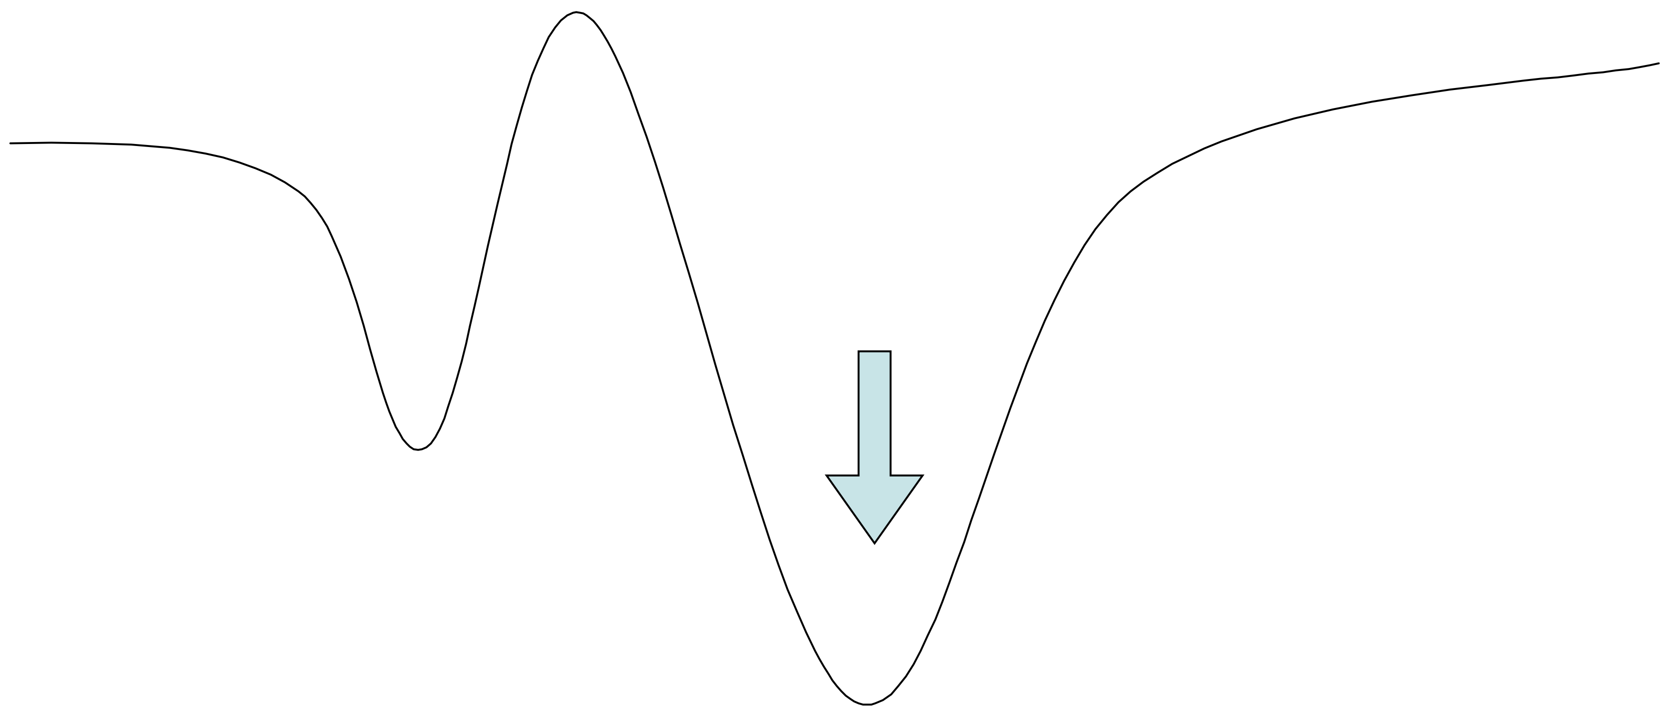

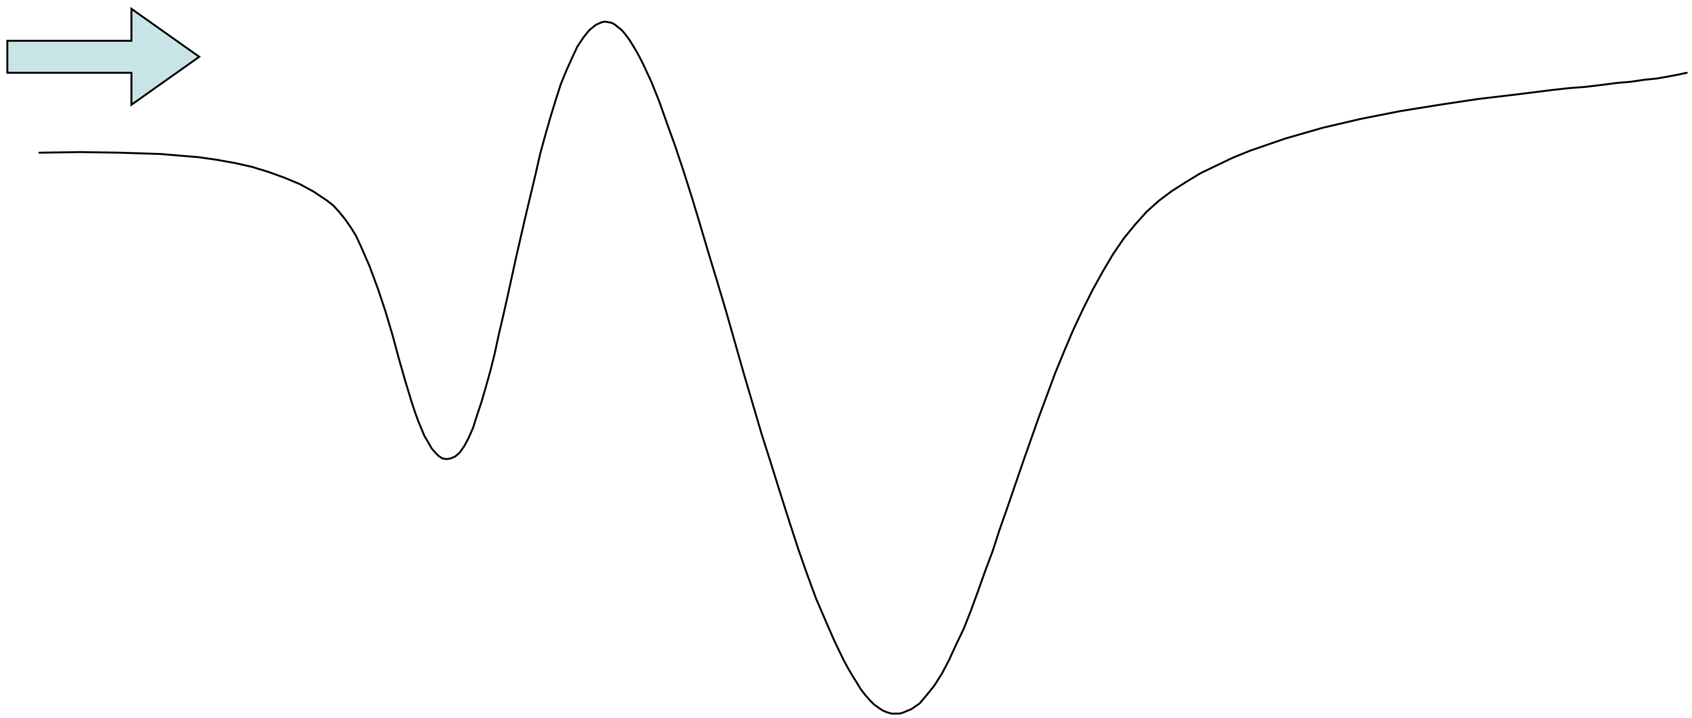

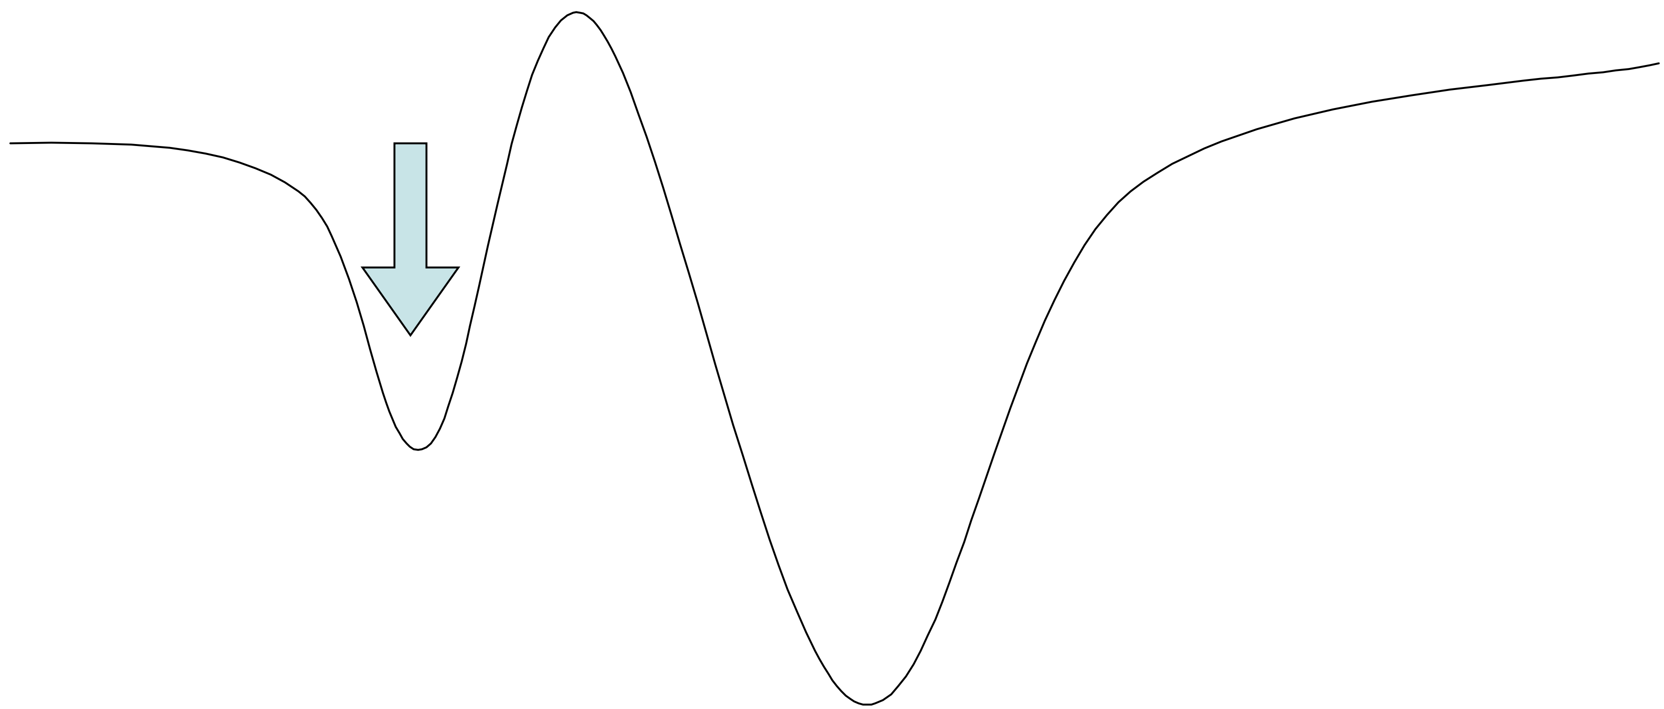

Finding the best-fit

Finding the best-fit means minimizing the statistic value. There are many algorithms available to do this in a computationally efficient fashion (see Numerical Recipes).

Most methods used to find the best-fit are local i.e. they use some information around the current parameters to guess a new set of parameters. All these methods are liable to get stuck in a local minimum. Watch out for this !

The more complicated your model and the more highly correlated the parameters then the more likely that the algorithm will not find the absolute best-fit.

Finding the best-fit II

Sometimes you can spot that you are stuck in a local minimum by using the XSPEC error or steppar commands. These both step through parameter values, error in the vicinity of the current best-fit and steppar over a user-defined grid, and thus can stumble across a better fit. Crude but sometimes effective.

You can do this in a semi-automated fashion by using a local minimization algorithm and following this with the error command with the ability to restart if a new minimum is found during the search.

Global Minimization

There are global minimization methods available - simulated annealing, genetic algorithms, … - but they require many function evaluations (so are slow) and are still not guaranteed to find the true minimum.

A new technique called Markov Chain Monte Carlo, which provides an intelligent sampling of parameter space, looks promising but it is not yet widely available (i.e. I’ve not added it to XSPEC - yet).

Dealing with background

o Unless you are looking at a bright point source with Chandra you will probably have a background component to the spectrum in addition to the source in which you are interested.

o You can include background in the model but this is complicated and is not usually used.

o The usual method is to extract a spectrum from another part of the image or another observation. Spectral fitting programs then use both the source and background spectra.

o If the background spectrum is extracted from a different sized region than the source then the background spectrum is scaled by the spectral fitting program (using the BACKSCAL keyword in the FITS file).

Spectra with few counts

o Be careful if you have few photons/bin. Chi-squared is biased in this case with fluctuations below the model having more weight than those above, causing the fit model to lie below the true model.

o A common solution is to bin up your spectrum so all the bins have > some number of photons. Don't do this - it loses information and introduces a bias that is difficult to quantify.

o Solutions are to use a different weighting scheme (I prefer the *weight churazov* option in XSPEC) or a maximum likelihood statistic (the "C statistic" - *stat cstat* in XSPEC).

o The problem with these options is that while they give best fit parameters they do not provide a goodness-of-fit measure.

Final Advice and Admonitions

- Remember that the purpose of spectral fitting is to attain understanding, not fill up tables of numbers.
- Don’t bin up your data - especially in a way that is dependent on the data values (eg group min 15).
- Don’t misuse the F-test.
- Try to test whether you really have found the best-fit.